



INSTITUT C.D. HOWE INSTITUTE

COMMENTARY

NO. 387

Predicting Recessions in Real-Time: Mining Google Trends and Electronic Payments Data for Clues

*Vast new sources of electronically recorded data, such as Google Trends, are both timely
and reflect the real-time intentions of millions (or billions) of agents.
Economists are studying their predictive power for economic conditions.*

Greg Tkacz

THE INSTITUTE'S COMMITMENT TO QUALITY

ABOUT THE AUTHOR

GREG TKACZ
is Associate Professor and
Chair of the Department of
Economics, St Francis
Xavier University.

C.D. Howe Institute publications undergo rigorous external review by academics and independent experts drawn from the public and private sectors.

The Institute's peer review process ensures the quality, integrity and objectivity of its policy research. The Institute will not publish any study that, in its view, fails to meet the standards of the review process. The Institute requires that its authors publicly disclose any actual or potential conflicts of interest of which they are aware.

In its mission to educate and foster debate on essential public policy issues, the C.D. Howe Institute provides nonpartisan policy advice to interested parties on a non-exclusive basis. The Institute will not endorse any political party, elected official, candidate for elected office, or interest group.

As a registered Canadian charity, the C.D. Howe Institute as a matter of course accepts donations from individuals, private and public organizations, charitable foundations and others, by way of general and project support. The Institute will not accept any donation that stipulates a predetermined result or policy stance or otherwise inhibits its independence, or that of its staff and authors, in pursuing scholarly activities or disseminating research results.

COMMENTARY No. 387
SEPTEMBER 2013
FINANCIAL SERVICES



Finn Poschmann
Vice-President, Research

\$12.00

ISBN 978-0-88806-907-8

ISSN 0824-8001 (print);

ISSN 1703-0765 (online)

THE STUDY IN BRIEF

Many official economic indicators are released with a time lag, released infrequently and often require revision. In this *Commentary*, I discuss new sources of electronically recorded data that are both timely and reflect the real-time intentions of millions (or billions) of agents. Specifically, I consider whether Google searches and the growth of electronic payments variables, such as debit and credit card transactions, would have predicted the 2008 – 2009 recession.

Not too long ago, Canadian empirical macroeconomic researchers would have to wait two months for the release of the monthly National Accounts in order to update their models and forecasts. However, in the last 10 years to 20 years, technological innovations have resulted in vast amounts of other data being recorded electronically and stored. New data series are now being generated at a rate faster than analysts can study them.

Due to the emergence of Google as the dominant search engine, its search-term usage can provide a snapshot of current group interests in numerous issues, such as economics, politics, health, etc. In principle, if many people are entering the same economic search terms, this could provide a clue about changing conditions, such as the onset of a recession.

I find that the usage of Google search terms “recession” and “jobs” could have predicted the recession up to three months in advance of its onset. However, since Google query data are only available from 2004, the time span studied in this *Commentary* is very short in the context of business cycles. Consequently, our study should be viewed as illustrative of the potential uses of electronic data. I also highlight the benefits and pitfalls that users of Google data may encounter in the context of economic monitoring.

C.D. Howe Institute Commentary® is a periodic analysis of, and commentary on, current public policy issues. Michael Benedict and James Fleming edited the manuscript; Yang Zhao prepared it for publication. As with all Institute publications, the views expressed here are those of the author and do not necessarily reflect the opinions of the Institute's members or Board of Directors. Quotation with appropriate credit is permissible.

To order this publication please contact: the C.D. Howe Institute, 67 Yonge St., Suite 300, Toronto, Ontario M5E 1J8. The full text of this publication is also available on the Institute's website at www.cdhowe.org.

Policymakers are clearly interested in recession forecasting since they would benefit from lead-time to plan and implement counter-cyclical policies that could potentially lessen recessionary impacts.

Currently, however, recession forecasting is complicated by the publication lags for many important macroeconomic variables, which makes difficult the real-time identification of an incipient recession.

For example, in the case of the most recent recession, Cross and Bergevin (2012) date its Canadian beginning to November 2008. However, in the middle of October, Canadian policymakers would have had access only to GDP data for August, so the recession recognition lag can be long. Relying on the optimistic August data, Prime Minister Harper on October 10, 2008 declared that, "This country will not go into recession next year and will lead the G7 countries."

Six weeks later, after receiving updated macroeconomic data, Harper had to change his tune. On November 23, he said, "The most recent private-sector forecasts suggest the strong possibility of a technical recession at the end of this year and beginning of the next."¹ Clearly, given the rapid speed at which the economy deteriorated in late 2008, having had access to more timely data could have allowed policymakers to recognize and respond to the recession more quickly.

Not too long ago, Canadian empirical macroeconomic researchers would have to wait two months for the release of the monthly National Accounts in order to update their models and

forecasts. However, in the last 10 to 20 years, technological innovations have resulted in vast amounts of other data being recorded electronically and stored. New data series are now being generated at a rate faster than analysts can study them.

For macroeconomic purposes, it would be useful to use electronic data that capture activity in a broad segment of the economy, are available on a timely basis, and can be measured with little or no error. In this context, payments innovations such as debit and credit cards satisfy many of these criteria. They are available, in principle, on a daily basis and, therefore, can track consumer spending virtually in real time. Because Canadian GDP data are available only after a two-month lag, researchers have been studying more readily available debit and credit card data for the purpose of predicting GDP before its release. Termed "nowcasting" (Galbraith and Tkacz 2013a), such data make possible the immediate assessment of extreme shocks to the economy, like the attacks of September 11, 2001.

Another new data source that has recently been receiving attention from economists is Internet search data. Due to the emergence of Google as the dominant search engine, its search term usage can provide a snapshot of current group interests in numerous issues, such as economics, politics, health, etc. In principle, if many people are entering the

The author gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada as well as the comments from several reviewers of this paper.

1 *Vancouver Sun*, November 26, 2008.

same economic search terms, this could provide a clue about changing conditions, such as the onset of a recession. For example, Wu and Brynjolfsson (2009) use Google data to predict housing prices, since searches on terms such as "real estate" or "housing prices" tend to be correlated with buyer or seller intentions. Google search data have also been successfully used for the production of labour market forecasts,² and their usefulness as leading economic indicators has been explored by some central banks.³

Given that there have been few studies involving Google Trends data for Canada, I discuss in the next section this source's peculiarities and provide examples of time series for search terms involving popular Canadian economic concepts. This exercise also highlights some of the problems involved when using these data, such as extreme observations that may be unrelated to the topic of interest, seasonality and significant revisions that occur as the time span lengthens. Therefore, unlike data that are compiled by Statistics Canada, Google Trends data require further filtering.

Next, I discuss predicting recessions using payments data, which include debit card, credit card and cheque payments. I also discuss a traditional recession indicator, the interest-rate yield spread, defined as the difference between a long-term government bond yield and a treasury bill yield. Atta-Mensah and Tkacz (2001) found this yield difference to be one of the more reliable recession indicators since the 1960s. This variable will be used as a benchmark against which the new indicators will be compared.

Following Atta-Mensah and Tkacz (2001), I present a model that allows users of these new data

sources to predict, in retrospect, the probability of a recession occurring between 0 and 24 months in advance. I found that a Google search series of the word "recession" spiked one month prior to the 2008/09 recession, which coincided with the onset of the US recession. For longer-run forecasts, defined as those 18 and 24 months in advance, I found that the yield spread remains the best leading indicator of a recession.

I must caution, however, that the empirical work in this *Commentary* deals only with this most recent business cycle, as Google Trends data are available only from January 2004. This means that our analysis is limited, containing but one recession, so I make little effort to defend the statistical significance, or lack thereof, of some findings. Nevertheless, I hope that the work presented here will help guide future use of Google Trends for macroeconomic monitoring.

In the final section, I summarize our findings and highlight the benefits and caveats for use of these data in policymaking. Although I uncover many problems, I also find that analysts now have a new data source that can potentially quantify what had previously been non-quantifiable, such as consumer confidence sentiments. However, much exploratory work remains to be done to find the proper search terms that best capture such emotions.

GOOGLE TRENDS DATA

This section begins by discussing how Google Trends data are constructed and then illustrates some of the issues that users may encounter through an examination of the search results of some common economic terms.⁴

2 See, for example, Ettredge, Gerdes and Karuga (2005); Choi and Varian (2009, 2012); Askatas and Zimmerman (2009) and D'Amuri and Marucci (2009).

3 See Suhoy (2009) at the Bank of Israel and McLaren and Shanbhogue (2011) at the Bank of England.

4 Google Trends data can be accessed at <http://www.google.ca/trends>.

Construction of the Data

Google Trends allows users to download a time series of a particular search term in a specific category in a specific geographic region over a specific period of time. As Choi and Varian (2012) note, the time series obtained are not actually the raw number of downloads, but rather index numbers that range in value from 0 to 100. Providing index numbers solves the problem of increase in Internet usage, where the volume of searches could simply be related to the growth of Internet users or, more importantly, the growth in the number of Internet devices that permit Google searches, such as smart phones. The date that corresponds to the highest number of searches for a particular term over a particular time period is assigned a value of 100. It must also be noted that the values downloaded, according to Choi and Varian (2012, p. 3), are "computed using a sampling method, and the results therefore vary a few per cent from day to day." This could be problematic for replicating research, since if data are downloaded on a different day, even during the same time span, the series might look slightly different.

Choosing Search Categories

Apart from isolating geographic areas, another useful feature of Google Trends is the ability to isolate a category under which the search is to be conducted. This can be very important if the search term can be interpreted in different ways according to the context.

For example, suppose an analyst interested in labour market trends chooses to study the search pattern for the word "jobs." And suppose further that the search is conducted over both "All Categories" and the "Jobs and Education" sub-category. Both series show a consistent pattern of decreasing interest in jobs during any calendar year, reaching a dramatic trough during Christmas week (see Figure 1). However, interest in jobs tends to spike during the first week after Christmas,

perhaps due to resolutions or renewed resolve. This information on its own could be useful to employers, as help-wanted ads are likely to generate more job applications in January than in December according to these series.

In terms of overall trends, I observed a decreasing-to-flattening trend for "job" searches from 2004 to 2008. But job interest began trending upward at the beginning of 2009 through early 2010 as the impact of the recession (see Figure 1) began to affect the labour market. From mid-2010 to early 2013, I observed a very slight downward trend, which would correspond to the post-recession labour market recovery.


The largest variance between the two job search series – All Categories and Jobs and Education – is the October 2011 spike that I observed in "All Categories." This corresponds to the death of Apple Inc. founder Steve Jobs on October 5, 2011, heightening an interest in "jobs" that was unrelated to labour market conditions. This serves as a cautionary example for users of such data, especially for economic purposes.

Although it would be tempting to automate the process of downloading a vast number of these series, feeding them into a forecasting model and then generating a single forecast for every period, analysts must first verify the potential causes of outliers. Otherwise, the October 2011 spike in "jobs" might have resulted in an unusually high unemployment rate forecast. At least in this example, the availability of a suitable sub-category could have prevented such an error from occurring.

Search for Some Common Economic Terms

Many terms in economics and finance are unique to the field. Consequently, one can perform broad searches of common terms using "All Categories." In Figure 2, I present All Categories search results for six terms, exclusively for Canada, over the January 2004 to January 2013 period. The terms: monetary policy, fiscal policy, inflation, GDP, exchange rate and unemployment rate.

Figure 1: Search Term “jobs” Canada – All Categories and Jobs & Education Category, January 2004 to January 2013, Weekly



Note: Highest search value scaled to equal 100.


Source: Google Trends data.

Monetary policy and fiscal policy both seemed to generate sporadic interest prior to 2008, with many periods being assigned values of zero due to lack of search volume. However, there appeared to be cyclical tendencies, which seemed to coincide with important policy announcements. In the case of monetary policy, these seemed to correspond with interest rate announcements, while for fiscal policy, many spikes appeared to occur at federal budget time. Both generated interest during the 2008/09

recession, but nothing out of the ordinary compared to other periods.

On the other hand, inflation searches appeared to be more volatile prior to the recession, when they seemed to exhibit a slight downward trend. Post-recession, the peaks and troughs of this series appeared to correspond to those of monetary policy, so interest in this topic apparently was driven by Bank of Canada fixed-action dates.

Figure 2: Search Terms “monetary policy,” “fiscal policy,” “inflation,” “gdp,” “exchange rate,” “unemployment rate” Canada – All Categories, January 2004 to January 2013, Weekly



Note: Highest search values scaled to equal 100.

Source: Google Trends data.

GDP, for its part, seemed to have a strong cyclical component, with troughs occurring between May 2008 and August 2008. Although this is mere speculation, it is possible that this drop can be attributed to college and university students being out of school, which might have produced fewer searches for academic purposes.

Meanwhile, exchange rate searches were dominated by a large surge in late 2007. This coincided with the period when the Canadian dollar reached an all-time high of about US\$1.10. Because of this extreme observation, it is difficult to detect any other regular occurrences or trends.

Finally, unemployment rate searches seemed to follow the known movements of the actual unemployment rate, as this series trended downward prior to the recession, then rose rapidly as the recession hit (see Figure 9). Thereafter, it gradually started falling again, much like the unemployment rate itself. Interestingly, this series, much like GDP, seemed to generate fewer queries during the summer months. The hypothesis that fewer macroeconomics students during the summer are behind this lull must again be forwarded.

In short, a visual inspection of these data shows that, apart perhaps from the unemployment rate, none of these series appear to display unusual movements prior to, or during, the recession. In fact, the unusual seasonality displayed by GDP and unemployment rate that seemed to follow the university academic year suggest that these may not be the best candidate recession predictors.

Some Possible Recession Search Terms

The number of search terms that one can consider for predicting recessions is limited only by one's imagination. However, I found potential value only in the obvious "recession," and "jobs," which Figure 1 shows jumped as the recession hit.

To understand whether policymakers could have used either of these terms to predict the recession in real-time, I first need to understand how search series for these two terms have evolved with the advent of new data. As discussed earlier, the historical series can be revised as a new maximum value is assigned the number 100. For jobs and recession, I considered what the series looked like for seven different time periods. All samples began in January 2004 and ended, respectively, in January 2007, 2008, 2009, 2010, 2011, 2012 and 2013.

For the search term recession in the time period ending in 2007, the series looks quite different from the other periods (see Figure 3). The reason is that recession was not a common search term prior to 2007. Canada had not experienced a

formal recession since 1992, and the economy was performing relatively well over this particular sample period, so few searches were done using this term.

By contrast, when I end the sample in January 2008, these earlier observations are superseded by a spike in December 2007. This coincides with the formal beginning of the US recession, so this is a situation where US reports made Canadians think and worry about recessions.


Extending the sample to January 2009 encompasses the start of the Canadian recession in November 2008. There are notable spikes in new recession searches, reflecting widespread media discussion about the recession and the federal government acknowledgement in late November 2008 that the economy might, indeed, be headed for a recession. These spikes lessened the impact of the December 2007 observation, which now is not as notable relative to the new observations.

One of the implications of such a large historical revision to this series from prior to the recession is that it would not have been a very reliable indicator in real time. Instead, I needed to observe the evolution of search results for this term during an actual recession prior to using it for practical forecasting purposes.

Finally, the samples ending in January 2010, 2011, 2012 and 2013 result in relatively minor revisions to the history of this series, which implies that, going forward, the "recession" series now has a sufficiently long and varied track record to be used without fear of major revisions, since I have observations that actually span a legitimate recession.

Unlike the recession series, the series for the jobs search term seems to not be greatly impacted by historical revisions, so this is a series that could have potentially been used in real time to help policymakers predict the 2008/09 recession (see Figure 4). Since I restricted myself to search results in the Jobs and Education category, I do not appear to have any unusual outliers in this series.

**Figure 3: Search Term "recession" Canada, All Categories,
Samples beginning in January 2004 and ending between January 2007 and January 2013, Weekly**




Note: Highest search values scaled to equal 100.

Source: Google Trends data.

What is most striking, however, is the regular seasonal variation, where searches on jobs every December are about one-half of those experienced every January. In fact, the search terms spike the week following Christmas, suggesting that many people postpone their job search during the holidays and resume immediately thereafter,

perhaps with renewed interest and hope. To use this series for forecasting purposes, I therefore need to remove the seasonality. This can be achieved using a standard seasonal adjustment procedure such as X12, or one can simply track the year-over-year growth rates, which is the practice I follow in my empirical application.

**Figure 4: Search Term "jobs" Canada, Jobs & Education Category,
Samples beginning in January 2004 and ending between January 2007 and January 2013, Weekly**



Note: Highest search values scaled to equal 100.

Source: Google Trends data.

OTHER DATA

Payments Data: Credit Cards, Debit Cards and Cheques

Through the cooperation of the Canadian Payments Association (CPA) and the Canadian Bankers Association (CBA), Galbraith and Tkacz

(2013a, 2013b) compile a payments data set that includes the total value and volume of credit card transactions, the total value and volume of debit card transactions and the total value and volume of small (<\$50,000) cheques that clear between banks.

When nowcasting GDP, Galbraith and Tkacz (2013a) focus on the changes to the average value

of each payment type. From 2004 to 2009, each debit transaction averages about \$45, each credit card transaction about \$110, and each cheque about \$1,100. In terms of year-over-year growth in their average value, these three types of payments all began to fall around October 2008 and recover shortly after May 2009 (see Figure 5). Consequently, these transactions could be correlated with the onset of a recession and, as with Google Trends data, they share the property of timeliness in the sense that these data can, in principle, be obtained on a daily basis. However, unlike Google Trends, one does not need to worry about historical revisions.

A TRADITIONAL RECESSION INDICATOR

Atta-Mensah and Tkacz (2001) forecast recessions using 20 different indicators, including interest rates, stock market indicators and monetary aggregates. The series that best predicts recessions at most horizons and over a long time span is the slope of the yield curve, defined as the difference between a long-term bond rate and a short-term interest rate. This variable usually has peak predictive power between 12 and 18 months in advance of a recession's onset, and the linkage to economic growth is normally thought to operate through monetary policy. If the short rate is low relative to the long rate, which can be viewed as a proxy for the equilibrium or policy-neutral rate, then monetary policy is thought to be expansionary and its impact on the economy will be felt 12 months to 18 months later. In the case where the short rate exceeds the long rate, such that the yield curve is inverted, then monetary policy is thought to be tight, so economic activity will slow in 12 months to 18 months.

Many past recessions have been preceded by an inversion of the yield curve, and I retain this variable in this *Commentary* as a solid benchmark against which the new electronic variables are

compared. The most recent recession was preceded by about 18 months by a relatively mild inversion of the yield curve, so it may yet have had some predictive ability for the 2008/09 recession (see Figure 6).


MODELS

Our objectives are to determine whether some of the variables defined above were coincident or leading indicators of the 2008/2009 recession. If they were leading indicators, then what is the horizon over which they led the recession. I also try to quantify the goodness-of-fit of each individual indicator. For this purpose, I use a so-called probit model that allows linking movements in these variables to the occurrence of the 2008/09 recession through probabilities (see box for more details).

Since the Google Trends data begin in 2004, our estimation period is very short in macroeconomic terms. Furthermore, our payments data end in 2010, so our estimation period spans only the six years between 2004 and 2009, which is barely one business cycle. Because of this short sample period, our results should be interpreted as illustrative rather than definitive and should serve only as a guide for analysts interested in monitoring these data going forward.

I also need to caution that the results I present are not generated out-of-sample; that is, I do not use data that would have been available to analysts in real time, so I do not precisely replicate the situations that would have been encountered by policymakers prior to the last recession. This is a constraint I cannot overcome since the lack of recessions between May 1992 and September 2008 meant that our dependent variable was simply a sequence of zeroes during this period. However, going forward, analysts can now use this model to generate real-time recession forecasts, as I have both recessionary and non-recessionary periods, if data from 2008/09 are included.

**Figure 5: Credit, Debit and Chequing Transactions
Value Per Transaction, Year-over-Year Growth, Monthly**



Sources: Canadian Payments Association, Canadian Bankers Association, Galbraith and Tkacz (2013a and b).

PREDICTING THE 2008/09 RECESSION

In Table 1, I present the estimated parameters for all six of our indicators at the eight different forecast horizons, for a total of 48 estimated equations. For each horizon, I highlight the indicator that generates the highest goodness-of-fit as captured by the pseudo- R^2 .

In Table, I also state the signs that I expect for the parameter β for each indicator. For the yield

spread, it should be negative: when the difference between the short-term rate and the long-term rate is high, monetary policy is considered to be expansionary and the probability of a recession should be low. For "recession" and "jobs" searches, I expect positive signs. If searches for recession increase, this suggests more people are worried about one actually occurring; for jobs, the higher the growth in the number of searches relative to the previous year, the higher the number of people either looking for work or worrying about their

Box 1: Probability of Being in a Recession – Technical Explanation of the Probit Model

Given our interest in the 2008/09 recession, I divide the economy into two possible states: It is either in a recession, or it is not. As such, our variable of interest is binary, since it can only take two values. Let $R_t = 1$ if the economy is in recession in period t , and 0 if it is not. From Cross and Bergevin (2012), R_t would equal 1 for November 2008 to May 2009 inclusively, and 0 otherwise.

Since I have a binary dependent variable and continuous indicators, I can link movements of the continuous indicators to the binary dependent variable through probabilities, which can be achieved via the following probit model:

$$(1) \quad P(R_t = 1) = F(\alpha + \beta X_{t-k})$$

Equation (1) models the probability of being in a recession at time t based on the value of an indicator X at time $t-k$, where k takes the values 0, 1, 2, 3, 6, 12, 18, 24. I estimate the model at a monthly frequency, so k represents the lags of X in months. I choose to estimate the model at a monthly frequency since the recession dates are available on a monthly basis, and because our credit card data are currently only available at this frequency. $F(\cdot)$ is the normal cumulative distribution function, so this allows us to map the continuous values of X_{t-k} in the range of 0 to 1. The unknown parameters α and β are estimated numerically.

The goodness-of-fit of (1) for each indicator at each horizon k is determined using the pseudo- R^2 statistic of Estrella (1998), which is analogous to the standard R^2 statistic in simple linear regression models. The pseudo- R^2 lies between 0 and 1, and the higher its value, the higher its goodness-of-fit. Intuitively, if a model predicts a recession with a probability of 1.0 in a period in which a recession actually occurs, and a probability of 0.0 in a period in which there is no recession, then the pseudo- R^2 would equal 1. Alternatively, if the model predicts a recession with a probability of 1.0 in all periods when a recession does not occur, and a probability of 0.0 in all periods in which a recession actually occurs, then the pseudo- R^2 would equal 0.


jobs and, therefore, the higher the probability of a recession.

For the payments variables, I expect the parameters to be negative. If average transaction values are growing, then this signals that households are spending more, so the likelihood of a recession should be lower.

The results from Table 1 suggest that the Google Trends data generate the best in-sample fit during short horizons; that is, from zero to three months. In other words, these six Table 1 scenarios would

have been the best predictors of the 2008/09 recession, if forecasts were generated in July, August, September and October 2008. The term recession, in particular, generates a high pseudo- R^2 one month ahead of the recession. Looking back to October 2008, the month following the failure of Lehman Brothers and the subsequent interventions by central banks around the world to pump liquidity into the financial system, these events could have driven more people to search for recession in Google.

Figure 6: Yield Spread 10-Year-and-Over Government Bond Rate minus 90-Day T-Bill Rate, Percent, January 1959 to January 2013, Monthly



Sources: CANSIM series V122487 (long-term rate) minus series V122484 (short-term rate). Recession dates from Cross and Bergevin (2012).


To better appreciate the predictive power of “recession,” in Figure 7 I plot the probabilities of a recession as predicted by our model using this variable. The probabilities are zero in most non-recession periods and are close to 1.0 during the actual recession. The only notable false signal generated by this indicator occurred in December 2007, when talk of a US recession likely was responsible for the relatively large number of Google searches in Canada.

The growth of searches for jobs is not as impressive a recession indicator as “recession,”

itself, but nevertheless it did generate many non-zero probabilities during the actual recession (see Figure 8). Its major drawback is that it continued generating non-zero probabilities even after the recession ended. This is not inconsistent with the observation that employment growth tends to lag GDP growth. It also suggests that this search term may be a better labour market than recession indicator.

In fact, the growth of jobs searches follows unemployment rate searches closely, with both series peaking after the recession (see Figure 9).

Figure 7: Recession Probabilities
Using "recession", 1-month forecast ($K = 1$)




Source: Author's calculations.

Given that labour force data is often survey-based and, therefore, subject to sampling variability, it is always useful for analysts to supplement their labour market toolboxes with additional indicators. The correlation between the actual unemployment rate and the growth rate of jobs searches is +0.63, but rises to +0.79 when growth of jobs is lagged six months, suggesting that this search term could be a useful leading indicator of the unemployment rate. Nevertheless, I leave the study of the usefulness of Google Trends for Canadian labour market data to future work.

The yield spread remains the best long-run recession predictor, being statistically significant (and its parameter being of the correct sign) at months 18 and 24 (see rows for $k=18$ or 24 in Table 1). The payments variables, meanwhile, behave like the Google Trends data in the sense that they are statistically significant (and of the correct sign) for short-run forecasting. The re-emergence of their statistical significance, notably for debit and cheques, at month 24 was somewhat unexpected and probably a statistical artifact, given the short sample period. Nevertheless, considering

Figure 8: Recession Probabilities
 Using Year-over-Year Growth of Searches for "jobs", coincident forecast ($K = 0$)



Source: Author's calculations.

that these are variables that are timely and can be observed with little to no delay, they remain valuable for the purpose of nowcasting economic activity.


DISCUSSION AND POLICY IMPLICATIONS

In this *Commentary*, I examined the potential usefulness of Google Trends for monitoring and forecasting Canadian economic conditions and, in the process, highlighted some of the advantages and

pitfalls of these data. I also illustrated how these and other recently compiled electronic payments data could have been used to predict the onset of the 2008/09 recession. My findings suggest that movements in Google Trends data were correlated with the recession and, in fact, could have been used to predict it up to three months in advance.

These new data open many new avenues for research, both theoretical and applied, and many of these will be of use to policymakers. Here are some benefits and caveats:

Figure 9: Unemployment Rate and Growth Rate of Searches for "jobs"
January 2005 to January 2013, Monthly



Source: Author's calculations.

Benefits

- (1) Google Trends data are very timely. They are aggregated on a weekly basis and, therefore, should be particularly well suited for nowcasting.
- (2) The data can be disaggregated by country, province and even city, when there are a sufficiently large number of observations. This can be especially useful for regional economic analysis, where the data may be less readily available.
- (3) If chosen judiciously, some search terms could help quantify emotions, such as love, fear, hate, etc., which have received relatively little attention in economics. This could help understand, for example, movements in Consumer Confidence Surveys.
- (4) In the short run, some of the most immediate benefits of Google Trends data appear to be in the analysis of labour market trends. Early positive results in this regard have already been noted for the United States, United Kingdom,

Table 1: Probit Model Results

	$P(R_t = 1) = F(\alpha + \beta X_{t-k})$					
	Variable (X)					
	Financial	Google Trends		Payments Variables		
	<i>Spread</i> (-)	<i>"recession"</i> (+)	<i>"jobs"</i> (+)	<i>Credit</i> (-)	<i>Debit</i> (-)	<i>Cheque</i> (-)
$K = 0$						
$\hat{\beta}$	0.65	0.08	0.11	-0.21	-0.59	-0.30
<i>t</i> -stat	2.84	3.64	3.05	-3.32	-3.42	-2.97
Pseudo- R^2	0.20	0.52	0.30	0.23	0.29	0.19
$K = 1$						
$\hat{\beta}$	0.58	0.28	0.10	-0.16	-0.42	-0.18
<i>t</i> -stat	2.68	1.20	3.13	-2.85	-2.96	-2.06
Pseudo- R^2	0.16	0.82	0.28	0.15	0.17	0.08
$K = 2$						
$\hat{\beta}$	0.51	0.14	0.09	-0.15	-0.32	-0.13
<i>t</i> -stat	2.46	2.08	3.12	-2.58	-2.38	-1.46
Pseudo- R^2	0.12	0.73	0.24	0.12	0.10	0.04
$K = 3$						
$\hat{\beta}$	0.43	0.06	0.08	-0.12	-0.23	-0.06
<i>t</i> -stat	2.13	3.85	2.96	-1.99	-1.75	-0.66
Pseudo- R^2	0.09	0.43	0.19	0.06	0.05	0.01
$K = 6$						
$\hat{\beta}$	0.23	0.03	0.07	-0.08	-0.04	-0.02
<i>t</i> -stat	1.15	2.37	2.59	-1.08	-0.28	-0.19
Pseudo- R^2	0.02	0.10	0.14	0.02	0.00	0.00

Table 1: CONTINUED

	$P(R_t = 1) = F(\alpha + \beta X_{t-k})$					
	Variable (X)					
	Financial	Google Trends		Payments Variables		
	<i>Spread</i> (-)	<i>[recession]</i> (+)	<i>[jobs]</i> (+)	<i>Credit</i> (-)	<i>Debit</i> (-)	<i>Cheque</i> (-)
$K = 12$						
$\hat{\beta}$	-0.25	0.03	0.05	0.03	-0.02	0.52
t-stat	-1.10	1.91	1.33	0.19	-0.09	2.50
Pseudo- R^2	0.02	0.06	0.04	0.00	0.00	0.18
$K = 18$						
$\hat{\beta}$	-1.21	0.01	-0.04	0.22	0.02	0.19
t-stat	-1.87	0.37	-0.64	0.97	0.07	1.34
Pseudo- R^2	0.20	0.00	0.01	0.02	0.00	0.03
$K = 24$						
$\hat{\beta}$	-5.68	0.12	0.15	0.04	-0.91	-0.69
t-stat	-2.41	1.25	1.85	0.20	-2.08	-2.95
Pseudo- R^2	0.48	0.03	0.10	0.00	0.12	0.25

Note:

Expected parameter signs are in parentheses. Highest Pseudo- R^2 for each horizon is highlighted. If the t-stat is greater than 1.985 in absolute terms then the parameter β should be considered to be statistically different from zero with 95% confidence. Spread is in levels; all other variables are measured as growth rates.

Germany and Israel by other authors, while in this *Commentary* the data suggest that this may also be the case for Canada.

- (5) These data could be incorporated into data-intensive forecasting models, such as factor models, and used to produce forecasts that can be verified against those generated by theory-driven models. They can also be used to help shape the

judgment of analysts who monitor short-run fluctuations in key economic variables.

Caveats

- (1) One needs to choose Google Trends search terms carefully, since some may have a meaning that is different from the intended economic meaning.

Steve "Jobs" is but one example. Consequently, analysts need to study outliers and think about potential non-economic causes for these observations.

- (2) Since the search data are available only from 2004 and since they are scaled from 0 to 100, some search term data may still be subject to relatively large historical revisions.
- (3) Google Trends data are compiled based on surveys of the number of search terms each period and, therefore, the downloaded series are subject to some sampling variability, even on a daily basis. This could make replicating some research results difficult if the underlying series is changing.
- (4) Search results data may be skewed by generations, as younger people tend to be more active on the Internet than seniors. This could explain, for example, the seasonal variations observed in some of the downloaded series, as lower search results of economics terms during the summer months may be related to the end of college and university classes.
- (5) Google does not provide the absolute number of search results for every period. This could be problematic if the underlying data are generated by relatively few search terms.
- (6) It may be difficult to incorporate some search terms into formal economic models. Most variables in models represent quantities or prices, and it is not always clear with some search terms if the variations are driven by quantities or prices. Therefore, one may not know if they should serve as proxies for quantities or prices.
- (7) Language needs to be taken into consideration. Quebecers, for example, are more likely to search for a term in French rather than English, so "jobs" might generate relatively few results in Quebec relative to "emplois."
- (8) Technology evolves rapidly, so while Google is the dominant search engine in 2013, it could be supplanted in the future. Consequently, users of

electronically generated data should always be aware of the impact of technological innovations that could affect the data being studied.

CONCLUSION

The nature of economic data available to analysts is rapidly evolving, as technological innovations allow for the storage, recording and retrieval of a vast amount of information at relatively low cost. This can potentially improve the quality and accuracy of economic monitoring in real time, which could allow policymakers to more quickly respond to deteriorations in economic conditions.

In this *Commentary*, I demonstrated the potential usefulness of new data sources for the specific purpose of predicting the 2008/09 recession. Given the short time span of some new data sources, this exercise should be viewed as illustrative rather than definitive, as it highlights some of the advantages and pitfalls that analysts may encounter when using them. As such, this *Commentary* is intended to serve as a starting point for new research going forward.

The primary new data source discussed is Internet query data, specifically the data that can be obtained from Google Trends. There are literally unlimited series that can be obtained from this source, as each downloaded series is unique for a given search term, geographic location and time span. I downloaded a handful of series for some relatively common search terms and found a generic term such as "recession" to be correlated with the onset of the 2008 recession, which preceded the release of official GDP data from Statistics Canada by at least two months. However, I leave it as a future exercise to locate the optimal combination of search words that could have predicted the recession even further in advance since, as I mentioned, the combination of search terms is limited only by one's imagination.

I also considered the usefulness of electronically recorded debit and credit card transactions, as well as cheques that clear through the payments system, as potential leading predictors of a recession.

Although these data are proprietary and difficult to obtain by the general public, their existence suggests that analysts could exploit such data to monitor economic conditions in real time.

Given that the analysis of these new timely, high-frequency indicators is in its infancy, I believe going forward that, as the time span expands and

the properties of these new sources are better understood, they will serve as useful inputs for macroeconomic monitoring. Their timeliness can potentially help policymakers recognize downturns as they occur and obtain a better measure of their severity in real time, thereby helping to devise appropriate policy responses.

REFERENCES

- Aruoba, S. B., F. X. Diebold and Chiara Scotti. 2009. "Real-Time Measurement of Business Conditions." *Journal of Business and Economic Statistics* 27:417-427. October.
- Askatas, Nikolaos, and K. F. Zimmerman. 2009. "Google Econometrics and Unemployment Forecasting." *Applied Economics Quarterly* 55:107-120. June.
- Atta-Mensah, Joseph, and Greg Tkacz. 2001. "Predicting Recessions and Booms Using Financial Variables." *Canadian Business Economics* 8(3):30-36. February.
- Choi, Hyunyoung, and Hal Varian. 2012. "Predicting the Present with Google Trends." *The Economic Record* 88:2-9. June.
- . 2009. "Predicting Initial Claims for Unemployment Benefits." Working Paper. Google Inc. July.
- Cross, Philip, and Philippe Bergevin. 2012. *Turning Points: Business Cycles in Canada Since 1926*. Commentary 366. Toronto: C. D. Howe Institute. October.
- D'Amuri, Francesco, and Juri Marucci. 2009. "Google It! Forecasting the US Unemployment Rate with a Google Job Search Index." Institute for Social and Economic Research. Working Paper No. 2009-32. November.
- Estrella, Arturo. 1998. "A New Measure of Fit for Equations with Dichotomous Dependent Variables." *Journal of Business and Economic Statistics* 16:198-205. April.
- Ettredge, Michael, John Gerdes and Gilbert Karuga. 2005. "Using Web-Based Search Data to Predict Macroeconomic Statistics." *Communications of the ACM* 48(11): 87-92. November.
- Galbraith, John W., and Greg Tkacz. 2013a. "Nowcasting GDP: Electronic Payments, Data Vintages and the Timing of Data Releases." Working Paper 2013s-25, CIRANO. Available at: http://www.cirano.qc.ca/publications_detail.php?lang=fr&id=2013s-25
- . 2013b. "Analyzing Economic Effects of September 11 and Other Extreme Events Using Debit and Payments System Data." *Canadian Public Policy* 39(1): 119-134. March.
- McLaren, Nick, and Rachana Shanbhogue. 2011. "Using Internet Search Data as Economic Indicators." *Bank of England Quarterly Bulletin* Q2, 134-140. June.
- Suhoy, Tanya. 2009. "Query Indices and a 2008 Downturn: Israeli Data." Bank of Israel, Discussion Paper No. 2009.06. July.
- Tkacz, Greg. 2010. "An Uncertain Past: Data Revisions and Monetary Policy in Canada." *Bank of Canada Review*, Spring, 41-51. March.
- Wu, Lynn, and Erik Brynjolfsson. 2009. "The Future of Prediction: How Google Searches Foreshadow Housing Prices and Sales." Working Paper, MIT Sloan School of Management. December.

NOTES:

NOTES:

NOTES:

RECENT C.D. HOWE INSTITUTE PUBLICATIONS

August 2013	Gros, Barry. <i>Ontario Pension Policy 2013: Key Challenges Ahead</i> . C.D. Howe Institute Commentary 386.
August 2013	Johnson, David. "Identifying Alberta's Best Schools." C.D. Howe Institute E-Brief.
August 2013	Daniel Schwanen. "Uneasy Birth: What Canadians Should Expect from a Canada-EU Trade Deal." C.D. Howe Institute E-Brief.
August 2013	Dachis, Benjamin. "How Ottawa Can Deliver A Reformed Canada Post." C.D. Howe Institute E-Brief.
July 2013	Boyer, Marcel, and Sébastien Boyer. "The Main Challenge of Our Times: A Population Growing Younger." C.D. Howe Institute E-Brief.
July 2013	Laurin, Alexandre, and Finn Poschmann. "Treading Water: The Impact of High METRs on Working Families in Canada." C.D. Howe Institute E-Brief.
July 2013	Dachis, Benjamin. <i>Cars, Congestion and Costs: A New Approach to Evaluating Government Infrastructure Investment</i> . C.D. Howe Institute Commentary 385.
July 2013	Busby, Colin, and Alexandre Laurin. "The 8 Percent Solution: A Sensible Tax Compromise for Albertans." C.D. Howe Institute E-Brief.
June 2013	Ragan, Christopher. "The Seductive Myth of Canada's "Overvalued" Dollar." C.D. Howe Institute E-Brief.
June 2013	Worswick, Christopher. "Improving Immigrant Selection: Further Changes Are Required Before Increasing Inflows." C.D. Howe Institute E-Brief.
June 2013	Morgan, Steven G., Jamie R. Daw, and Michael R. Law. <i>Rethinking Pharmacare in Canada</i> . C.D. Howe Institute Commentary 384.
June 2013	Howitt, Peter. <i>From Curiosity to Wealth Creation: How University Research Can Boost Economic Growth</i> . C.D. Howe Institute Commentary 383.
May 2013	Beaudry, Paul, and Philippe Bergevin. "The New 'Normal' for Interest Rates in Canada: The Implications of Long-Term Shifts in Global Saving and Investment." C.D. Howe Institute E-Brief.

SUPPORT THE INSTITUTE

For more information on supporting the C.D. Howe Institute's vital policy work, through charitable giving or membership, please go to www.cdhowe.org or call 416-865-1904. Learn more about the Institute's activities and how to make a donation at the same time. You will receive a tax receipt for your gift.

A REPUTATION FOR INDEPENDENT, NONPARTISAN RESEARCH

The C.D. Howe Institute's reputation for independent, reasoned and relevant public policy research of the highest quality is its chief asset, and underpins the credibility and effectiveness of its work. Independence and nonpartisanship are core Institute values that inform its approach to research, guide the actions of its professional staff and limit the types of financial contributions that the Institute will accept.

For our full Independence and Nonpartisanship Policy go to www.cdhowe.org.



C.D. HOWE
INSTITUTE

67 Yonge Street, Suite 300,
Toronto, Ontario
M5E 1J8